Report for NASA AISRP Grant NNG05GA94G

Project Title:
A Neural Map View of Planetary Spectral Images for Precision Data Mining and Rapid Resource Identification

Erzsébet Merényi, PI, Rice University
Robert H. Brown Co-I, University of Arizona, Tucson, AZ
William H. Farrand, Co-I, Space Science Institute, Boulder, CO
Thomas Villmann, Collaborator, Klinik für Psychotherapie, Universität Lepzig, Germany
Colin Fyfe, Collaborator, University of Paisley, Scotland

Project period: 11/1/2004 – 10/31/2007 plus extension year 11/1/2007 – 10/31/2008
Project URL: http://www.ece.rice.edu/∼erzsebet/HYPEREYE.html

Contents

1 Project Summary (from original proposal)

This project follows up on a current three-year AISRP grant, NAG5-10432, which will end 8/2004. It addresses a pressing need for rapid yet intelligent analysis of voluminous multi- and hyperspectral images in order to extract key data and generate knowledge. Spectral imaging plays a leading role in remote identification of surface materials of Earth (Landsat, AVIRIS, Hyperion), Mars (Pathfinder, MGS, MER), the Jovian system (Galileo NIMS), the Saturnian system (Cassini VIMS) and other solar system bodies. Hyperspectral sensors, in particular, enable detailed identification through the complexity of signatures measured in hundreds of narrow spectral bands. The challenge in automated and fast (real time, on-board) interpretation of these huge images calls for massively parallel algorithms, as well as it requires sophisticated algorithms for optimal knowledge extraction. Properly utilized, Artificial Neural Networks (ANNs) can provide both.

The current project engineered ANNs, specifically Self-Organizing Maps (SOMs) and their hybrids for efficient and sophisticated clustering and classification of spectral images, developing custom modules supported by commercially available components. Based on some of the latest theoretical research on SOMs, the tools we developed, jointly with European experts, are powerful for distinguishing a large number of spectral classes and for the discovery of "interesting" but uncommon and spatially very small classes. We use information theoretically principled SOM approaches, which increases power and confidence in autonomous data mining. We demonstrated the effectiveness and high quality of data analysis on sample IMP spectral images, Cassini VIMS Jupiter fly-by imagery, AVIRIS and other data representing typical challenges in NASA's missions. We propose to advance these computational intelligence capabilities in three ways: 1) We will add significantly new theoretical strength to information extraction modules. 2) The software, HyperEye, will be made transferable to other users through high-level graphic interface, augmented software design, tutorials and wrapping, opening an important phase of technology infusion that will take recent and future developments into the user community. 3) We will directly participate, using our methods and software, in analyses of spectral images forthcoming from the Mars Exploration Rovers and Cassini VIMS Saturn orbital tour, and (pending its funding) a Pluto/icy satellites spectral analysis project.

The 'neural' core of our software is already suitable for implementation in high-speed massively parallel hardware (which could be an on-board analysis capability), as it was one of the original objectives of our work. We are pursuing that line of development outside of this project proposal and, if successful, we anticipate using the hardware to support this work as well.

This project is a collaboration between computer science and space science investigators at Rice University, University of Arizona, and the Space Science Institute, Boulder, CO.

2 Accomplishments in Year 3 plus no-cost extension year 4, 11/1/06 - 10/31/08

2.1 Completion of project tasks (algorithm and software development, data analysis)

This project suffered a loss of programming help at the end of year 2, in part due to the increase in competitive salaries in the Houston area, which our budget was unable to match. After almost a year of unsuccessfully trying to hire a replacement, we asked and received a one-year no-cost extension. We also restructured efforts to maximize the output of this project for relevant science return. We dedicated more graduate student time to the project and focused more on scientific algorithm development and data analyses (Tasks 2–4) and somewhat less on the refinements of software augmentation (Task 1). This report summarizes year 3 and the extension year. Figures 1 and 2 show, respectively, a conceptual overview, and a top level layout of the functional components of our software environment, HyperEye. There are three facets of this environment: scientific algorithm development, software development, and data analysis, as reflected in the Tasks. Since HyperEye is focused on neural self-organized learning of high-dimensional manifolds, in order to produce detailed and precise segmentation and classification of highly structured data the modules we developed are specific to those needs, and the overwhelming majority of them are original algorithms (or implementations with original modifications) by the PI's group.

Task 1: Work on user interfaces, data handling and other support layers The overwhelming majority of this Task was completed in years 1–2, as previously reported. As we anticipated at the end of the previous period, during Year 3 (and 4) the need for expansion of data formats/handling was in the way of customized pre- and post-processing capabilities (data_digest, kappa-stat, Wilcoxon-rank in Figure 2, and other summarization scripts not shown in Figure 2) as scientific results produced by neural processing became more abundant and richer with new knowledge extraction modules created in this period (see Task 2). These new modules add to capabilities of previously developed ones such as clstat (class / cluster statistics summary), vecplot (a plotting package that has the intelligence to arrange plots such that all graphs show appropriately (for example, properly offset for viewing), class labels or legends are placed where they should be, etc.), and specter (an interactive image cube exploration tool).

Our previously reported meta-data facilities (that partly fall under data_digest) for the effective handling of ascii type input data (as opposed to image data that has spatial context) were further tested through a collaborative project with the Baylor College of Medicine's Cardiology Department, in which we have been analyzing clinical data.

Our core neural modules were updated to use QT. (The QT library is not indicated explicitly in Figure 2 as it is underlying the development components shown in the two top left columns.) Remodeling with QT a) serves to prepare our software for easy porting to multiple platforms in order to facilitate technology

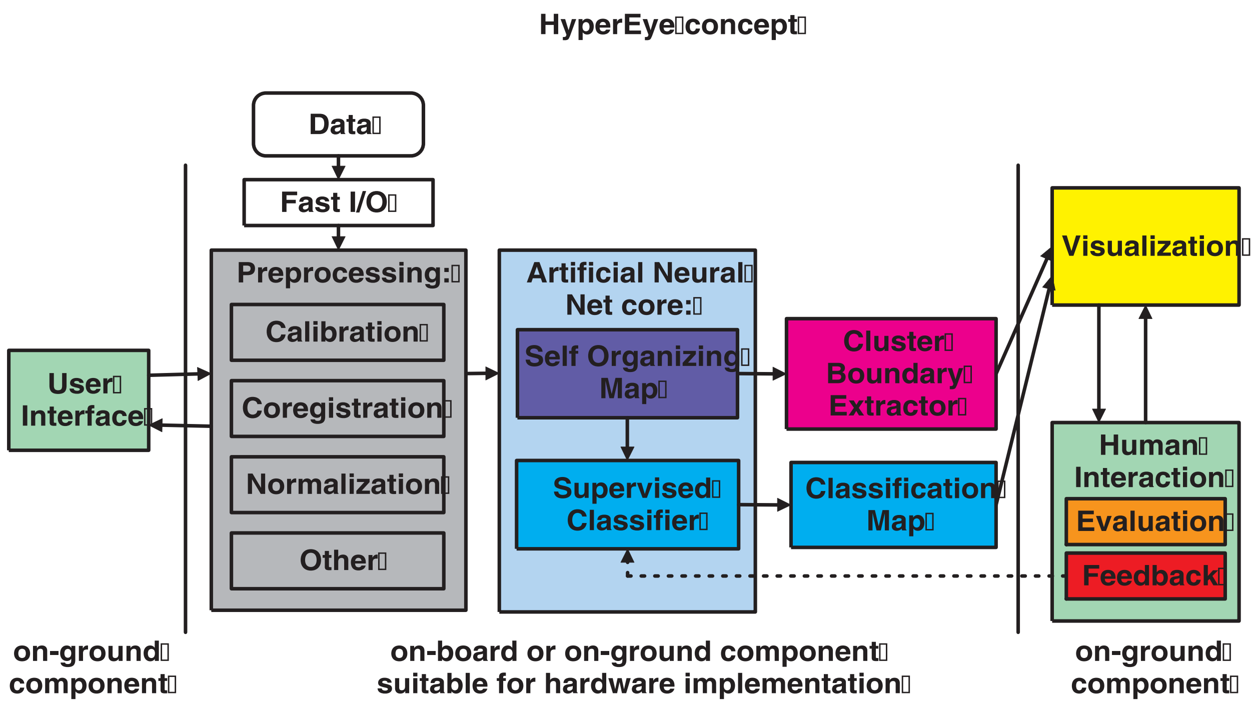

Figure 1: The conceptual structure of our software environment. A center idea is that the massively parallel neural algorithms can be implemented in hardware chips and embedded in autonomous, on-board, real-time processing and decision making systems. This project focuses on the development of smart neural clustering and classification, suitable for science information extraction from complicated and rich planetary spectral data.

transfer; (the actual porting to multiple platforms is beyond this project), b) is providing a support layer for updating/extending various existing functionalities. The QT library's visualization components aid work in Task 3 and elsewhere, assist meta-data parsing, and exporting of results to various graphic formats. It will also aid history recording of interactive user input during module operation. We have a visualization library (Vismod), built on QT (and also underlying the top level components), which provides a unified frame for displaying all data that we handle, including neural net components / layers, as well as user data, and has the hooks for using meta data for customized annotations. Development of customized visualization through plug-ins (slave modules to main modules) makes recompilation of a main module unnecessary when new capabilities are added to the plug-in, and a module can receive services from multiple plug-ins. This facilitates fast development of a variety of customized visualizations that can be used on demand by multiple modules. Our algorithm developments under Tasks 2 and 3 were greatly accelerated by the use of plug-ins.

Web demonstration of our core neural tools is at http://www.ece.rice.edu/∼erzsebet/HYPEREYE.html, complete with online documentation, demo data sets and tutorial.

Task 2: Advancement of ANN/SOM scientific knowledge extraction algorithms:

We wrapped up "relevance learning" that we had previously reported on. Then graduate student Major Michael Mendenhall (PhD August, 2006) worked with the PI on Generalized Relevance Learning Vector Quantization (following Hammer and Villmann, 2002; Villmann *et al.*, 2003) to assess the relative importances of hyperspectral data dimensions, which is a metric learning approach based on a supervised version of SOMs. To briefly recap, we investigated the GRLVQ because we have not seen good feature extraction for hyperspectral data, *i.e.*, schemes that preserved the distinction among the many classes that

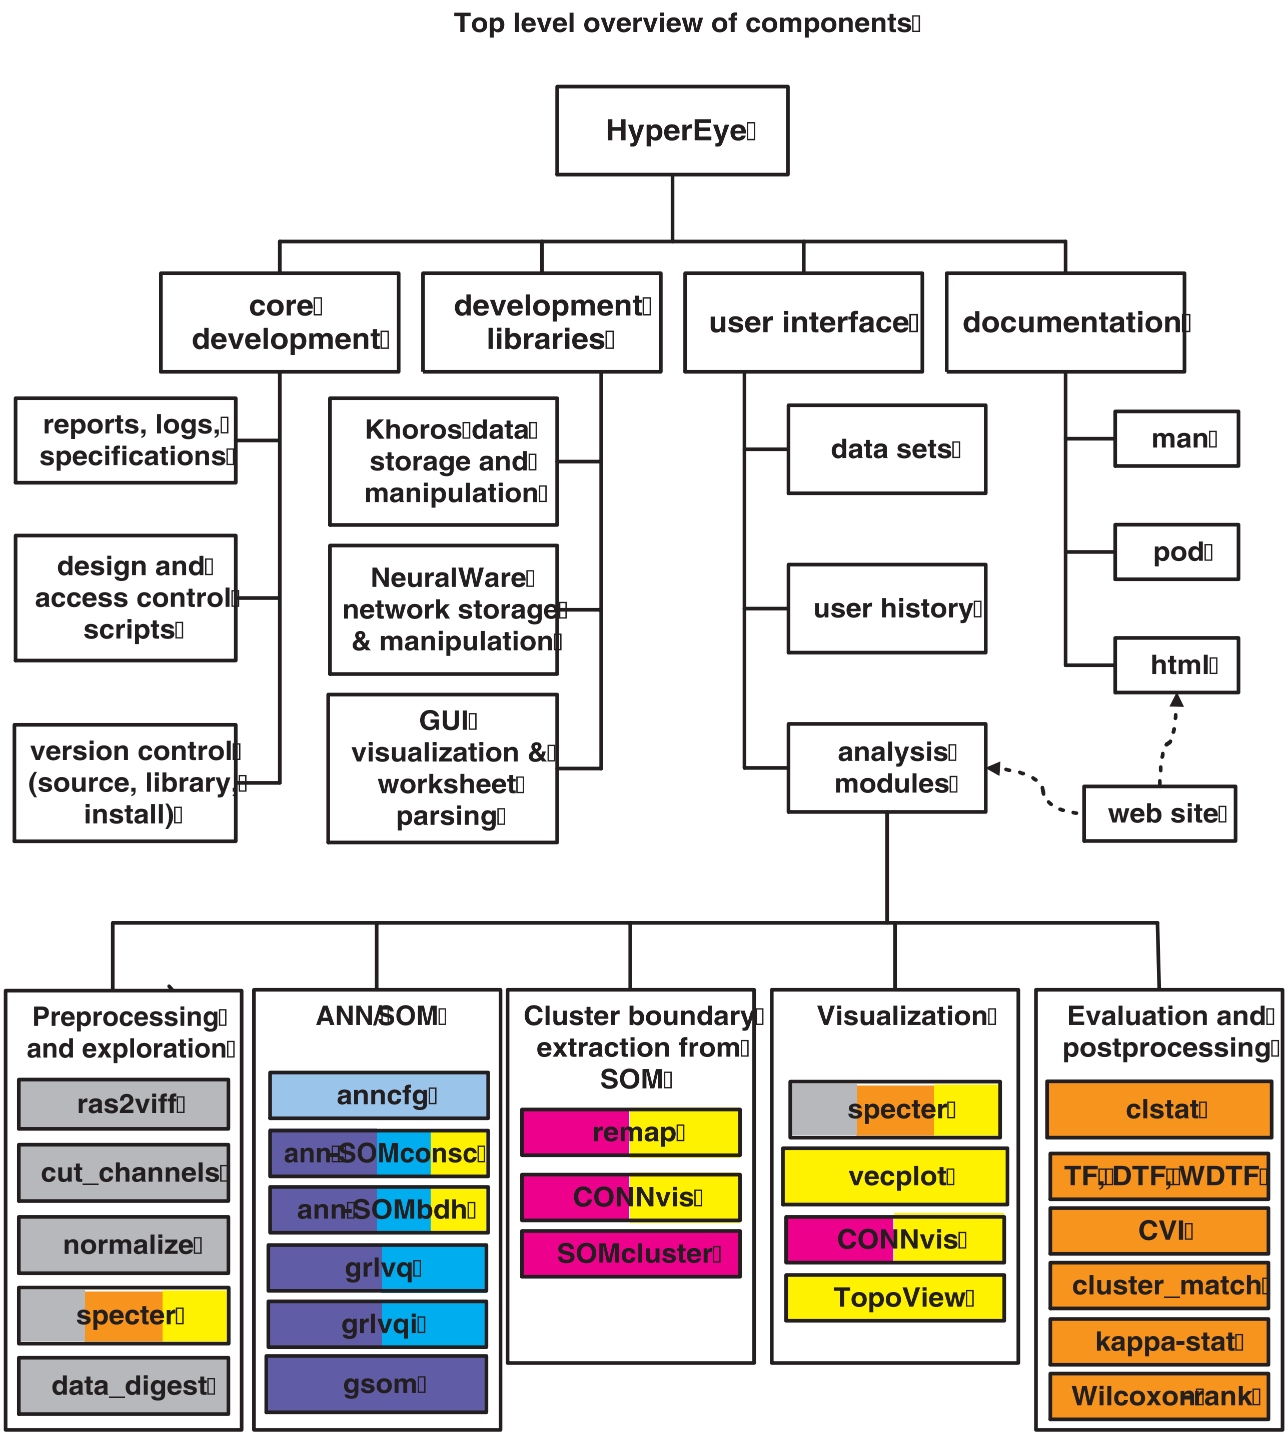

Figure 2: The main components of our software environment. The two topmost columns on the left (core development and development libraries) outline the supporting software components for the scientific algorithm development. Underlying the development libraries are our own visualization library (Vismod), and QT, not shown in this chart. The main scientific analysis engines are under "analysis modules". The modules' color codes indicate which functional groups they belong in Figure 1. Miscellaneous utilities (which are also modules) are not shown here for space constraints.

hyperspectral data is meant to distinguish. The GRLVQ is the first non-parametric method we know of, to measure the quality of the feature selection by the classification accuracy on the same data. Non-parametric approach is important for hyperspectral data where data models and priors are usually unavailable. We found an instability in the Hammer and Villmann scheme, devised and implemented a modified learning (GRLVQ-Improved, or GRLVQI) that stabilized the process and increased the classification accuracy and speed (grlvq and grlvqi modules in Figure 2). A *fundamental finding* of this research is that when GRLVQ is applied in the wavelet domain the wavelet coefficients that are selected as the N most relevant ones for a given classification, are not the same as the N largest magnitude coefficients! While this may be surprising, and is certainly different from the prevailing practice of wavelet coefficient selection for best modeling of data, it can be explained: The N *largest wavelet coefficients* ensure *minimum distortion* restoration of the data from the quantized encoding. In contrast, the coefficients with the N *largest relevances* provide *maximum information* relative to the class discrimination requirements of the given task. We achieve higher classification accuracy with a substantially reduced set of input features (in both the original reflectance data domain and in the wavelet transform domain), than with the full input. This could not be achieved with traditional feature extraction approaches such as PCA, or conventional selection of wavelet coefficients.

Following earlier conference papers (Mendenhall and Merényi, 2006a, 2006c), we completed and published a journal paper in IEEE Trans. Neural Networks (Mendenhall and Merényi, 2007).

We also put the finishing touch on our formerly reported investigation on SOM magnification control. This involved the implementation of the theory by Bauer, Der and Herrmann (1996), (hence our nickname, BDH, for the algorithm), in our ann-SOMbdh module, and its evaluation on data that the theory does not support but which data (such as hyperspectral imagery) may benefit the most from controlled SOM magnification for discoveries. The BDH algorithm was evaluated against the conscience learning SOM (ann-SOMconsc) for controled maximum entropy learning (as opposed to heuristics in ann-SOMconsc), and for the capability of exaggerating the representation of small clusters (whose existence is not known a priori), and thus increasing their detectability. We showed that at these two regimes the BDH worked well for "forbidden" data and published our findings with supporting case studies in the IEEE Trans. on Neural Networks (Merényi, Jain, Villmann, 2007).

As noted in our previous report, matching (reconciling) clusters across multiple images is a difficult task, which presents itself in many large data analysis projects. The problem for us arose with the Imager for Mars Pathfinder's Superpan octants, several of which we clustered. Since clusters are assigned labels arbitrarily within each separate clustering, it is hard to tell which cluster in image 1 describes the same surface cover as a given cluster in image 2. The PI's newest graduate student, Brian Bue, developed a cluster matching algorithm as part of his PhD qualifying semester project, and demonstrated initial success on synthetic spectral images. This algorithm, "cluster-match" in Figure 2, is now a module in our HyperEye collection, and will be further developed. A relevant related topic is the use of appropriate metrics in clustering or classification for particular classes of data. Spectral and hyperspectral data belong to the class of "Functional data". A metric recommended for spectral data is the Sobolev metric (Villmann et al., 2008). In the cluster matching algorithm, Brian implemented both Euclidean and Sobolev metrics for comparison and made a preliminary observation that using Sobolev metric resulted in better matching of clusters. Brian's contributions are at no cost to this project because he was awarded a NASA Graduate Student Research Fellowship. The goals of his fellowship are closely aligned with this project's goals, for mutual benefits.

Graduate student Kadim Tasdemir defended his PhD in April, 2008. His thesis work took a three-pronged approach to precise clustering of high-dimensional data. One component is the SOM visualization work under Task 3, another is an automated clustering of the SOM prototypes (SOMcluster module) which was created based on long-standing experiences with our interactive module for clustering the SOM (remap) and on experience with our new module CONNvis (see Task 3). The third is an original cluster validity index, CONN_Index, to assess the quality of a clustering without prior information. CONN_Index, a similarity measure based on data topology, has been shown to provide more useful measure for complicated structures than existing other cluster validity indices (Tasdemir and Merényi, 2007). CONN_Index can also be used for the evaluation of any prototype based clustering, not only for SOM prototypes. CONN_Index is implemented, along with several popular contender indices in our module CVI. We are in the process of submitting a journal paper "A validity index for complex cluster structures" to IEEE Pattern Recognition.

Our German collaborator, Dr. Villmann visited in June, 2008, during which time, as well as at Dagstuhl 2007 (March 2007), and ESANN 2008 in April 2008 (see activities in Section 2.2), we had intense discussions on further directions with relevance learning, metrics for functional data, and fuzzy labeled SOMs. Dr. Villmann also served on Major Mendenhall's PhD committee.

Task 3: Improvement of visualization and human interaction

Taşdemir and the PI proposed an *original visualization* of data manifolds. This representation is a weighted version of the so-called (binary) Adjacency Matrix produced by standard Delaunay triangulation. In our cersion a weighting of the connection between two Voronoi centroids is assigned, as the number of data points for which one centroid is the closest, the other is the second closest prototype. Hence the elements of this weighted adjacency matrix, which we call *Connectivity Matrix (CONN)*, express not only a binary connectedness of the manifold, but also the local density, i.e., how strongly various regions are connected. This greatly facilitates cluster capture in a noisy data set, including the identification of outliers, based on thresholds automatically determined from the statistics of the weighted Delaunay graph. Draped over an SOM, data of any dimensionality can be visualized with this representation, in 2 dimensions, contrary to existing data visualizations which are limited to up to 3-dimensional data. After a conference paper (Taşdemir and Merényi, 2006), we now have an accepted journal paper in press (Taşdemir and Merényi, 2008a). Kadim implemented this visualization algorithm as a HyperEye plug-in module, CONNvis, with interactive user-driven query features.

CONNvis is highly complementary to other topology representation and evaluation tools that graduate student Lily Zhang created earlier (a plug-in, TopoView, for flexible visualization and evaluation of topology violations, and the functions TF, DTF and WDTF, based on Villmann *et al.*, (1997) and Zhang and Merényi (2006)), all of which are envisioned as components of a future unified tool to monitor the quality of SOM mapping during learning.

Task 4: Application of HyperEye algorithms to real scientific data Further scientific interpretation was done on clusterings and classifications that were produced in the previous cycle from IMP SuperPan spectral images. Ours is the first comprehensive classification of the SuperPan images acquired in 1997, due in part to calibration difficulties that introduced a "mosaic" effect within octants. The comprehensive classification allowed to examine a large statistics of the compositional variations, which revealed some differences with previously reported compositional trends. Following preliminary conference papers (Farrand *et al.*, 2005 and Wright *et al.*, 2005), we published these findings in the International Mars Journal (Farrand *et al.*, 2008).

We performed clustering on a spectral image from the MER Spirit. These data are very similar to the IMP data but much cleaner and without the mosaic effect. Thanks to this, and to new capabilities such as CONNvis, this analysis was much faster than that of the IMP Superpan octants. This clustering shows great consistency with previous analysis done with independent methods, as well as increased compositional detail and discovery of a previously unreported feature. We published this in an invited paper at the SPIE Defense and Security Symposium, Space Exploration Technologies, and in a paper at the Discovery Science 2008 conference (Merényi, Taşdemir , Farrand (2008) and Taşdemir and Merényi (2008b)).

This project also supported joint work with AISRP project NNG05GA63G (PI Eliot Young), for which we have been developing neural classifier models for the prediction of surface temperature and grain size of surfaces of distant planetary bodies (such as Pluto) from synthetic ice spectra. Two internal presentations were given at Rice Quantum Institute Symposia by graduate student Lily Zhang. Preliminary results show that both temperature and grain size can be predicted by our models to satisfactory degree of accuracy (70-80% and 90-100% of the data points are predicted within 5% error for temperature and grain size, respectively). We are presently completing a noise sensitivity study to assess more realistic scenarios, and expect to prepare a publication in the near future.

Work on terrestrial analogs with well known ground truth is a very important component of our development of computational intelligence algorithms for planetary applications. Analysis on an AVIRIS terrestrial hyperspectral image of an urban scene with a large number of varied cover types (man-made as well as natural materials), groups of which show subtle differences (such as concretes or asphalts of various ages). In addition, several unique objects cover a few pixels only, which poses a challenge to clustering algorithms to discover those extremely small clusters. Our analysis of this scene shows such discoveries, and an excellent match of the signatures of extracted clusters with published field spectra of the respective surface types, for example, green tennis court, red tile roof, and more (Merényi, Csathó and Taşdemir , 2007).

We used terrestrial hyperspectral imagery to support a data compression study in a joint AISRP project (Tamal Bose, PI). The compressions schemes developed by PI Bose and his student Bei Xie are applied to the hyperspectral image, and subsequently classified with our neural tools to test if the desired class distinctions have been retained. Feedback from the classification results govern the refinement of the compression algorithms. We use neural classification for this purpose because we do not know of other classifiers that perform as well as our hybrid SOM-ANN classifier, with 200-dimensional feature vectors and 20–30 classes. We reported this work in Xie *et al.*, (2007, 2008).

We were planning, as a new component, to create one or two "reality based synthetic" spectral images, which would combine the characteristic real properties of existing, large and complicated spectral images with the convenience of a completely labeled data set. This has been a sorely missing component in standard national repositories, which do not represent the sophistication of planetary spectral data and thus do not pose appropriate challenge and testbed for our algorithms. We did not need to do this work, however, because we learned about a group (DIRSIG) at the Rochester Institute of Technology, which has already accomplished this task, on a much more professional and sophisticated level than we would be able to achieve. The PI received a very complex and realistic sample hyperspectral data from RIT (approx. 200 image bands, the scene containing over 70 different surface cover classes that range, in size, from 1 pixel to tens of thousands of pixels). With our tools we produced a clustering that is a faithful mapping of the ground truth (Merényi, Tasdemir, Farrand, 2008). We also used this data to demonstrate the theoretical aspects of our algorithms in a recently submitted, solicited book chapter (Merényi, Tasdemir, Zhang, 2008). The PI is in conversation with the DIRSIG group about further possible developments to their synthetic hyperspectral imagery.

2.2 Publications, presentations, other related activities

Refereed journal and refereed conference proceedings

Most papers are downloadable at http://www.ece.rice.edu/∼erzsebet/publist-Merenyi.pdf

Erzsébet Merényi, Kadim Tasdemir, Lili Zhang, (2008), Learning highly structured manifolds: harnessing the power of SOMs, Chapter In "Similarity based clustering", Lecture Notes in Computer Science (Eds. M. Biehl, B. Hammer, M. Verleysen, T. Villmann), Springer-Verlag, submitted.

Tasdemir, K, and Merényi, E. (2008a), Exploiting the Data Topology in Visualizing and Clustering of Self-Organizing Maps, IEEE Trans. Neural Networks, in press.

Bea Csathó, Justin Rich, Erzsébet Merényi, Lynn Everett, Brian Bue, John Kimble and Chien-Lu Ping, (2008), Characterizing polar landscapes from hyperspectral imagery, Proc. Ninth Intl Conference On Permafrost (NICOP 2008) (Eds. D. L. Kane and L. M. Hinkel), Fairbanks, AL, June 27 – July 1, 2008.

Farrand, W. H., Merényi, E., Bell, J. III, Johnson, J., Murchie, S. and Barnouin-Jha, O (2008), Class maps of the Mars Pathfinder landing site derived from the IMP SuperPan: Trends in rock distribution, coatings and far field layering, The International Journal of Mars Science and Exploration http://www.marsjournal.org/ , Mars, 4:33–55 doi:10.1555/mars.2008.0004, July 11, 2008..

Tasdemir, K, and Merényi, E. (2008b), Cluster analysis in remote sensing spectral imagery through graph representation and advanced SOM visualization, Proc. 11th Intl Conf. on Discovery Science, DS-2008, Budapest, Hungary, 13–16 October, 2008, in press.

Villmann, T., Merényi, E.and U. Seiffert, (2008), Machine Learning Approaches and Pattern Recognition for Spectral Data, Proc. 16th European Symposium on Artificial Neural Networks, ESANN'2008, Bruges, Belgium, 23–25 April, 2008. pp. 433–444 (tutorial paper for special session.).

Merényi, E., K. Tasdemir, and W.H. Farrand (2008), Intelligent Information Extraction to Aid Science Decision Making in Autonomous Space Exploration, Proceedings of DSS08 SPIE Defense and Security Symposium, Space Exploration Technologies, March 17–18, 2008, Orlando, FL.(Ed. W. Fink), 6960, 6:9600M, (Invited) http://scitation.aip.org/dbt/dbt.jsp?KEY=PSISDG&Volume=6960&Issue=1.

Mendenhall, M.J., Merényi, E. (2008), Relevance-based Feature Extraction for Hyperspectral Images, IEEE Trans. Neural Networks., 19(4):658–672.

B. Xie, Tamal Bose and Merényi, E., (2008), Novel algorithms for optimal compression using classification metrics, Proc. IEEE Aerospace Conference, March 2008.

B. Xie, Tamal Bose, and Merényi, E., (2007), New Algorithms for the Classification and Compression of Hyperspectral Images, Proc. NASA Science and Technology Conference, College Park, Maryland, June 19 - 21, 2007.

Tasdemir, K. and Merényi, E., (2007), A new cluster validity index for prototype based clustering algorithms based on inter- and intra-cluster density, Proc. Intl Joint Conf. on Neural Networks (IJCNN 2007), Orlando, FL, August 12–17, 2007. IEEE Catalog number 07CH37922C

Merényi, E., Farrand, W. H., Brown, R. H., Villmann, Th,. Fyfe, C., (2007), Information extraction and knowledge discovery from high-dimensional and high-volume complex data sets through precision manifold learning, Proc. NASA Science Technology Conference (NSTC2007), College Park, Maryland, June 19 – 21, 2007. 11pp. ISBN 0-9785223-2-X

Erzsébet Merényi, (2008), Biologically inspired computation for intelligent autonomous exploration, SPIE Newsroom. http://spie.org/x2434.xml?parentid=x2418&parentname=Astronomy&highlight=x2418

Michael Biehl, Erzsébet Merényi and Fabrice Rossi (2007), Advances in computational intelligence and learning, Neurocomputing, 70(7-9):1117–1119.

Merényi, E., L. Zhang, and K. Tasdemir, (2007), Min(d)ing the small details: discovery of critical knowledge through precision manifold learning and application to on-board decision support, Proc. IEEE Intl Conference on Systems of Systems Engineering (IEEE SoSE 2007), San Antonio, TX, April 16–18, 2007.

Merényi, E., B. Csathó, and Tasdemir, K., (2007), Knowledge discovery in urban environments from fused multi-dimensional imagery, Proc. 4th IEEE GRSS/ISPRS Joint Workshop on Remote Sensing and Data Fusion over Urban Areas (URBAN 2007), Paris, France, April 11–13, 2007 (invited paper). EEE Catalog number 07EX1577

Merényi, E., Jain, A., Villmann, Th. (2007), Explicit Magnification Control of Self-Organizing Maps for "Forbidden Data", IEEE Trans. Neural Networks, 18(3): 786–797.

Recent honors, PI

• Member of Editorial Advisory Board, International Journal of Intelligent Computing in Medical Science and Image Processing (IC-MED Journal), from 2006.

• Associate Editor, Neurocomputing, 2006 - 2007, for special issue on "Advances in computational intelligence and learning"

• "By Invitation only" participation in Dagstuhl Seminar "Similarity-based Clustering", Int'l Conference and Research Center for Computer Science, March 25–30, 2007, Schloss Dagstuhl, Wadern, Germany.

International conferences program committees the PI served on

• IASTED International Conference on Computational Biology and Bioinformatics (CBB 2008), Orlando, Florida, USA, November 16–18, 2008.

• Seventh International Conference on Machine Learning and Applications (ICMLA08), Special Session on Application of Machine Learning in Constructing Biopatterns and Analyzing Bioprofiles, December 11–13, 2008, San Diego, California, USA

• 16th European Symposium on Artificial Neural Networks, ESANN'2008, Bruges, Belgium, April 23–25, 2008

• 15th European Symposium on Artificial Neural Networks, ESANN'2007, Bruges, Belgium, April 26–28, 2007

• Third International Conference on Intelligent Computing and Information Systems, ICICIS 2007, Cairo, Egypt, March 15–18, 2007

Invited presentations by the PI

• Toward autonomous on-board science: self-organized neural learning of highly structured manifolds. University of Paderborn, Heinz Nixdorf Institute, April 29, 2008.

• Intelligent Information Extraction to Aid Science Decision Making in Autonomous Space Exploration, DSS08 SPIE Defense and Security Symposium, Space Exploration Technologies, Orlando, FL, March 18, 2008. (With Kadim Tasdemir and William H. Farrand)

• Information extraction and knowledge discovery from high-dimensional and high-volume complex data sets through precision manifold learning, NASA Science Technology Conference (NSTC2007), College Park, Maryland, June 19 – 21, 2007. (Merényi, E., Farrand, W. H., Brown, R. H., Villmann, Th,. Fyfe, C.)

• Knowledge discovery in urban environments from fused multi-dimensional imagery, 4th IEEE GRSS/ISPRS Joint Workshop on Remote Sensing and Data Fusion over Urban Areas (URBAN 2007), Paris, France, April 11-13, 2007. (With B. Csathó and K. Tasdemir)

• SOM and GRLVQ in Remote Sensing Image Analysis, Dagstuhl Seminar on Similarity Based Clustering. Supported in part by the German Informatics Society. Schloss Dagstuhl Computer Science Center, Germany, March 25–20, 2007.

3 Remaining Work, Technology Transfer, Future Directions

Due to the slow down in years 3–4 for reasons explained in Section 2.1, we have a small amount of work, and a proportionally small amount of funds (approximately 2.5% of the total funding, in the order of $15K) remaining. We are in the process of applying for another no-cost extension, to be able to complete outstanding items. These fall into Tasks 1 and 4, outlined below.

Task 1, Work on user interfaces, data handling and other support layers (∼70 – 75% effort)

With the comprehensive ascii and meta data format in place we can implement more visualization functionalities that use meta data. This will be the single most useful addition to existing software. We have been using the ascii / meta data facilities to process planetary ice spectra for a joint project (AISRP NNG05GA63G, PI Eliot Young) and for a heart failure analysis jointly with the Baylor College of Medicine's Cardiology Department.

The update of HyperEye modules to use QT and our custom software support components (e.g., Vismod) will continue. History recording of interactive user actions during module operation may be added to these updated modules as priorities permit. We will also improve the error message system. Our plug-ins and some of our new modules are using the Vismod library. Attempt to update older, major modules with Vismod/QT, namely our data exploration module "specter", however, presented some hard surprises that necessitate analysis of the components' interaction, and re-evaluation of this approach for the particular module.

Both user and devlopment documentation will be updated and checked to reflect the current state of software at the end of this period.

Our website for public demonstration will be updated. Our current interactive web demo is easy to view from unix machines, but it requires a cumbersome communications set up from PCs. We plan to add a video tour of the core capabilities to the interactive demo.

Task 4, Application of HyperEye to real scientific data will be ∼ 25 – 30% of the effort. Terrestrial hyperspectral data will be analyzed to more comprehensively evaluate GRLVQ(I). The PI and Major Mendenhall plan to publish one more paper on this subject. We continue to work on Eliot Young's planetary ice spectra (in a joint AISRP project), which is now coming to fruition and we expect to write a paper in the next few months. Similarly, we continue to support Tamal Bose's compression project by classifying compressed / reconstructed hyperspectral data (as in Task 4 under Accomplishments).

We have not received Cassini VIMS data so far, for reasons of priority of other pressing urgencies for the VIMS PI (R. Brown) and team. We will do VIMS analysis if we receive data.

Technology transfer is technically possible, but non-trivial. The following issues are involved:

1) Interested collaborators can request, install and run our software "as is", on compatible Sun/unix platforms, and purchase the necessary third party components (khoros and NeuralWare commercial software) elements of which are utilized not only as libraries in our compiled executables but also in command lines of our pre- and post-processing scripts. We have no resources, however, to provide bug fixes or help for porting to other platforms. As noted in section 2.1 under Task 1, this grant project aimed at preparing HyperEye for platform independence but not at actual porting (which would exceed the level of support this grant provides).

2) Our TRL level is about 3-4. While we try hard to keep the software well organized, and maintain version control, some documents or scripts may need updating at any one time since the software is in constant development. (For example, we assign to all modules symbolic names by which they can be invoked from the unix command line. This assignment is done through automated scripts, but these scripts need update when a new module is introduced or when we decide to change the name of a module.) Therefore, at the time of a request there may be some discrepancies, and the user needs to be familiar with the general philosophy and schemes in order to reconcile these, or they may need to wait for us to do it.

3) More importantly than point 2), the use of our software assumes quite intimate familiarity with SOMs and ANNs, including the theories that we implemented in our neural knowledge extraction modules. The implementation of those theories / functionalities are not found in commonly available neural packages, and may not be described elsewhere than in our published papers. As we mentioned part of these are original

inventions by us, or original implementations and modifications from theories. It is the user's responsibility to become knowledgeable about the particular algorithms.

4) In addition to 3), using our core modules involves a substantial learning curve. The user manual for operating our ann-SOMconsc module, for example, is more than 20 pages, as can be seen at our demo web site. Several others have similarly heavy documentation, with Vecplot's over 50 pages. Because (as seen from 3) and 4)) successful data analysis with our modules requires qualified users, the most fruitful infusion / utilization of our capabilities has been through personalized services, in which we participate in the analysis of our collaborators' data (as described in section 2.1 under Task 4). Substantial departure from this model where, for example, a commercial-strength on-line help system can aid in more intuitive learning of the usage, where more robust trouble shooting as well as problem fixing services can be offered, would require raising the TRL level (which in turn would require funding commensurate with such work).

5) We can, however, train a limited number of collaborating colleagues in the use of our software on a participatory basis (i.e., shifting the neural analyses to them while working on a joint project), which appears a sensible next step for infusing this technology into the planetary community. This would follow, among others, the example of the developers of spectral mixture analysis in the 1990's (John Adams and his school at the University of Washington), where they extended such personal coaching along with software to small groups at a time. Through that activity the community gradually became familiar with both conceptual issues and techniques, as well as with practices of driving algorithms and interpreting results, while producing valuable science.

Future directions: From its conception, the neural processing in this project aimed at two goals: smart data analysis, and massive parallelism. The latter allows — in principle — implementation in massively parallel hardware that can be embedded in on-board or other autonomous processing and decision making systems. (Hardware development is not part of this grant project.) Hardware that can accomplish the large scale processing that our SOM-based algorithms represent has not yet been built. However, in the last few years the PI has been in touch with a group that has the interest as well as the capability and experience to accomplish this. As of summer, 2008, a bread-board model has been put to preliminary test which indicated that a factor of 1,000 to 10,000 gain in speed (compared to sequential Sun wokstations) can be achieved. This translates to the potential of processing a standard AVIRIS hyperspectral image cube with an SOM in a few seconds, thus nearing the ability of real-time on-board processing or sifting through Earth-based archives fast. While there is a long way from the bread-board model to an embedded functional chip, we will be pursuing this hardware implementation if we can find suitable support.

4 References (other than listed under section 2.2)

Bauer, H.U., Der. R., Herrmann, M. (1996), Controlling the Magnification Factor of Self-Organizing Feature Maps, Neural Computation, 8:757–771.

Farrand, W. H., **Merényi, E.**, Murchie, S., Barnouin-Jha, O. (2005), Spectral Class Distinctions Observed in the MPF IMP SuperPan Using a Self-Organizing Map., Proc. 36th Lunar and Planetary Science Conference, Houston, Texas, March, 2005. , (Extended abstract).

Hammer, B., Villmann, Th. (2002), Generalized relevance learning vector quantization, Neural Networks, 15:10593/41068.

Mendenhall, M.J., **Merényi, E.**, (2006c), Relevance-based Feature Extraction from Hyperspectral Images in the Complex Wavelet Domain, Proc. IEEE Mountain Workshop on Adaptive and Learning Systems (SMCals/06), Logan, Utah, July 24 – 26, 2006. pp. 24–29

Mendenhall, M.J., **Merényi, E.**, (2006a), Generalized Relevance Learning Vector Quantization for Classification Driven Feature Extraction from Hyperspectral Data, Proc. ASPRS 2006 Annual Conference and Technology Exhibition, Reno, NV, May 1-5, 2006.

Taşdemir, K., **Merényi, E.**, (2005), Considering Topology in the Clustering of Self-Organizing Maps., Proc. 5th Workshop On Self-Organizing Maps (WSOM 2005), 5 - 8 September, 2005, Paris, France. Accepted. pp. 439-446

Villmann, Th., Herrmann, R.Der, and Martinetz, Th. (1997), Topology Preservation in Self-Organizing Feature Maps: Exact Definition and Measurement, IEEE Trans. on Neural Networks, 8(2):256–266.

Villmann, T.,**Merényi, E.**, Hammer, B. (2003), Neural Maps in Remote Sensing Image Analysis, Neural Networks, Special Issue on Neural Networks for Analysis of Complex Scientific Data, 16:(3–4):389-403.

Zhang, L. and **Merényi, E.**, (2006), Weighted Differential Topographic Function: A Refinement of the Topographic Function, Proc. 14th European Symposium on Artificial Neural Networks, ESANN'2006, Bruges, Belgium, 26-28 April, 2006. pp. 13–18

5 Appendix: Acronyms

ANN	Artificial Neural Network
AVIRIS	Airborne Visible and Infrared Imaging Spectrometer, of NASA, JPL
BDH	algorithm for SOM magnification control by Bauer, Der and Herrmann (1996)
DTF	Differential Topographic Function (a topology preservation measure)
CONN	Connectivity Matrix
GSOM	Growing Self-Organizing Map, an ANN paradigm
GRLVQ	Generalized Relevance Learning Vector Quantization
GRLVQI	Generalized Relevance Learning Vector Quantization Improved
IMP	The Imager for Mars Pathfinder
MER	Mars Exploration Rovers
PCA	Principal Components Analysis
SOM	Self-Organizing Map, a neural network paradigm
TF	Topographic Function (a topology preservation measure)
UA	University of Arizona
VIMS	Visible-Infrared Mapping Spectrometer, Cassini mission
WDTF	Weighted Differential Topographic Function (a topology preservation measure)